BBC RD 1975/3

1969 THE QUEEN'S AWARD TO INDUSTRY
1974 THE QUEEN'S AWARD TO INDUSTRY

RESEARCH DEPARTMENT REPORT

Video digital filter study I - first order recursive form

J. O. Drewery, M.A., Ph.D.

Research Department, Engineering Division
THE BRITISH BROADCASTING CORPORATION January 1975

BBC RD 1975/3
UDC 621.372.54

VIDEO DIGITAL FILTER STUDY I – FIRST ORDER RECURSIVE FORM

J.O. Drewery, M.A., Ph.D.

Summary

This Report describes the theory and practice of using a particular kind of digital filter to process video signals. This filter behaves like a R-C integrator but its effect on the picture can be associated with various types of resolution, namely horizontal resolution, vertical resolution or even temporal. Combinations are also discussed. The theory of the filter is developed and amplitude/frequency characteristics, pulse response, pole plot and quantisation effects are discussed. Practical results are also outlined. The application of the filter will be discussed further in companion reports.

Issued under the authority of

P. Rainger

Head of Research Department

Research Department, Engineering Division,
BRITISH BROADCASTING CORPORATION

January 1975

(PH-134)

VIDEO DIGITAL FILTER STUDY I – FIRST ORDER RECURSIVE FORM

Section Title Page

Summary .... Title Page

1. Introduction 1

2. Background 1

2.1. The difference equation and basic elements 1
2.2. The z-transform 2
2.3. Poles and zeros 3
2.4. Elemental forms of filter 3
2.5. Quantisation effects 4

3. The first-order recursive filter – theoretical 6

3.1. The pole plot 6
3.2. The frequency characteristic 6
3.3. The pulse response 6
3.4. The relationship with the analogue integrator 6
3.5. The effects of quantisation 7

4. The design of a first-order recursive filter 9

4.1. The sampling frequency 9
4.2. The number of bits 9
4.3. The need for ECL 11
4.4. Circuit details 11
4.5. Observed performance 12

5. The vertical first-order recursive filter – theoretical 12

5.1. The relationship with the horizontal filter 12
5.2. The pole plot 14
5.3. The frequency characteristic 14
5.4. The pulse response 15

6. The design of a vertical first-order recursive filter 15

6.1. Circuit details 15
6.2. Observed performance 15

7. The combination of horizontal and vertical first-order recursive filters 15

8. Conclusions 18

9. References 20

© BBC 2004. All rights reserved. Except as provided below, no part of this document may be reproduced in any material form (including photocopying or storing it in any medium by electronic means) without the prior written permission of BBC Research & Development except in accordance with the provisions of the (UK) Copyright, Designs and Patents Act 1988.

The BBC grants permission to individuals and organisations to make copies of the entire document (including this copyright notice) for their own internal use. No copies of this document may be published, distributed or made available to third parties whether by paper, electronic or other means without the BBC's prior written permission. Where necessary, third parties should be directed to the relevant page on BBC's website at http://www.bbc.co.uk/rd/pubs/ for a copy of this document.

VIDEO DIGITAL FILTER STUDY I – FIRST ORDER RECURSIVE FORM

J.O. Drewery, M.A., Ph.D.

1. Introduction

The idea of a filter which operates on a signal specified only in sampled form is not new. In effect it is only one of a more general class in which data is manipulated in discrete form, and which has existed for centuries; the commonest examples are interpolation, prediction and smoothing. What is relatively new is the ability to do the computations on a wideband signal, such as that involved in broadcast television, in a very short time so that the required results can be produced as fast as the data is generated. If the data consists of samples from an analogue signal the newly computed data represents samples of a modified signal which can be recovered in its continuous form, provided that the sampling rate is sufficiently high. As the computations are carried out in 'real time', the device which performs them can then be called a filter, performing a function similar to that of its analogue counterpart. Such a device is sometimes called a digital filter although strictly speaking a name such as 'sampled-data filter' should be used, the former term being reserved for devices which operate on samples which have been first quantised and turned into digital form.

In the field of television a filter broadly satisfying the above definition has existed in one form for some time. As television is a sampling system in the vertical direction the well-known vertical-aperture corrector can be considered as a sampled-data filter because it produces modified line signals, i.e. vertical samples, using information from adjacent line signals. In fact it belongs to the class of transversal filters defined below.

Interest in the digital filter has been reawakened by the fact that with the advent of pulse-code modulation the audio or video signal must itself be sampled and quantised. It would clearly be inconvenient to turn the samples into a continuous function to carry out a filtering operation and re-sample the result. Moreover it would be unnecessary because any signal processing operation can, in principle, be expressed as the manipulation of sample values with a precisely defined mathematical relationship of output to input.

The great attraction of digital filters is this ability to form precisely defined outputs in response to defined inputs. Thus the problem of long-term stability does not exist. The drawback is that filter characteristics which are easy to realise using analogue techniques may be much harder to realise digitally.

It is not surprising that the properties of digital filters have been investigated with the aid of digital computers. The mathematical relationships between sample values can be set up as a programme which can be executed with any set of values as the input sequence, the resulting set of values being the output sequence. In this way the action of a digital filter can be simulated without the investigator's needing to build it physically. The only difference is that the filter does not operate in 'real time'.

The difficulty of this simulation approach is that it is not easy subjectively to appraise the distortions introduced by the filter. There is thus a pressing need for computers to be able to operate on sample sequences in 'real time'. But the sample rate necessary even for broadcast-quality sound signals is too high for existing computer peripheral equipment. Thus specially designed equipment or 'dedicated hardware' must be built whenever subjective appraisal is required. Needless to say real time video processing by computer is, at present, completely out of the question.

The subject of digital filters is growing rapidly and there already exists a formidable body of relevant literature. Most of it has been concerned with applications to sound-signal processing and so has appeared in the journals appropriate to that medium.[1] However, interest in video processing is also gaining momentum. This Report can do no more than outline the background of the subject and give some practical examples. For a general review the reader is referred to a recent IEEE publication.[2] The background material in the following section follows closely the treatment of Gold and Rader.[3]

2. Background

2.1. The difference equation and basic elements

Digital filters arise out of the need to express a computed output sample as a linear combination of input samples and previous output samples. The theory of digital filters is based on linear difference equations just as the theory of continuous analogue filters is based on linear differential equations. Such an equation relates the present value of the output, y, to, say, the m previous values of y and, say, the n previous values of the input, x. If the input and output samples occur at times kT where k is an integer then the difference equation, of order m, takes the form

$$B_0\, y(kT) = \sum_{i=0}^{n} A_i x(kT - iT) - \sum_{i=1}^{m} B_i y(kT - iT) \quad (1)$$

where A_i and B_i are constants. The basic elements of a filter for performing such a calculation are therefore a means of access to successive samples together with adders and multipliers. If the samples are consecutive in time the means of access to previous samples is naturally provided by delays of value T or, where digital arithmetic is used, by storage elements clocked at the sampling times. If the

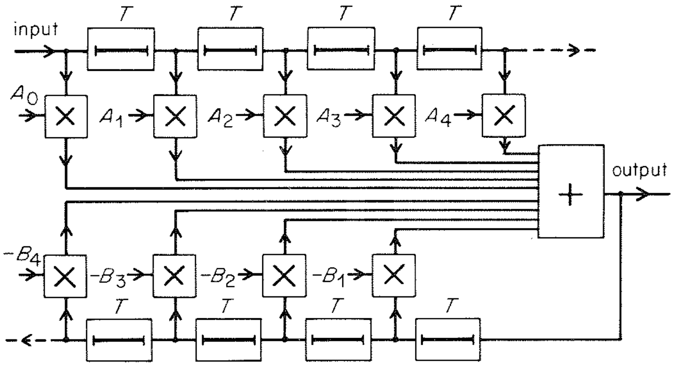

Fig. 1 - Direct-form realisation of general linear difference equation

operation of the filter can be expressed in the form of Equation (1), where B_0 equals unity by convention, it can be realised directly by the arrangement of Fig. 1. (Unfortunately, the specification of a digital filter is rarely expressed initially in difference equation form.)

As can be seen from Fig. 1 each output sample is computed from two distinct chains of samples corresponding to previous input and output samples. If the constants A_i are all zero except A_0 the output is derived only from previous output samples and the present input sample, and the arrangement is said to be purely recursive. If the constants B_i are all zero except B_0 the output is derived only from previous input samples and the arrangement is described as non-recursive. Alternatively it may be called transversal because of the resemblance to its analogue counterpart. Recursive and non-recursive filters are also termed infinite impulse response (IIR) and finite impulse response (FIR) filters respectively because of their behaviour when fed with the input sequence 1, 0, 0, ... which, in sampled systems, is analogous to the impulse in continuous systems.

2.2. The z-transform

Although the difference equation format of Equation (1) could be considered as the basic description of a digital filter it is not very helpful in predicting its response to generalised inputs, just as the linear differential equation is not very helpful for continuous systems. However, there is a mathematical procedure called the z-transform which, in sampled systems, is comparable with the Laplace transform used in continuous systems.

The analysis of linear continuous systems is based on the assumption that the input and output signals can be expressed as superpositions of complex frequencies of the form e^{pt}. The response of the system, if linear, is the superposition of the responses to the individual components and the system has thereby a transfer characteristic which is a function of the complex variable p. This analysis can be done with sampled systems but the representation is confused by the appearance of higher order components in the p-plane. Thus it is better to make a change of variable which takes this repetition into account. It is for this reason that the variable z is introduced.

The z-transform of a sequence $\{x_k\}$ whose members occur at discrete intervals kT, $k = 0, 1, 2, \ldots\ldots$ and are zero for $k<0$ is defined as

$$X(z) = \sum_{k=0}^{\infty} x(kT)z^{-k}$$

If Equation (1) is re-arranged into the form

$$\sum_{i=0}^{m} B_i y(kT - iT) = \sum_{i=0}^{n} A_i x(kT - iT)$$

and the z-transform applied to both sides, the result can be written

$$Y(z)\sum_{i=0}^{m} B_i z^{-i} = X(z)\sum_{i=0}^{n} A_i z^{-i} \qquad (2)$$

or

$$Y(z) = H(z)X(z) \qquad (3)$$

where $H(z)$ is the system transfer function expressed in terms of the z-transform. From Equation (2) $H(z)$ may be written in the form

$$H(z) = \sum_{i=0}^{n} A_i z^{-i} \Bigg/ \sum_{i=0}^{m} B_i z^{-i} \qquad (4)$$

which is the ratio of two rational polynomials in the variable z^{-1}. The response to any input may now be determined using Equation (3) by multiplying the z-transform of the input by the z-transform of the system and taking the inverse transform.

If the input samples are taken from a general complex exponential so that

$$x(kT) = e^{pkT}$$

the output samples are also samples of the same complex exponential multiplied by a quantity which is independent of k because of the special property of the exponential. That is

$$y(kT) = F(e^{pT})e^{pkT}$$

where $F(e^{pT})$ is the system transfer characteristic. If these expressions are substituted into the linear difference equation the result is that

$$F(e^{pT}) = \sum_{i=0}^{n} A_i e^{-ipT} \Bigg/ \sum_{i=0}^{m} B_i e^{-ipT} = H(e^{pT})$$

Thus the variable z is associated with the Laplacian variable p through the equation

$$z = e^{pT} \tag{5}$$

where T is the sampling period. In particular the response characteristic for a sampled sinusoidal input where $p = j\omega$ is given by

$$H(e^{j\omega T}) = \sum_{i=0}^{n} A_i e^{-ji\omega T} \Bigg/ \sum_{i=0}^{m} B_i e^{-ji\omega T} \tag{6}$$

Examination of Equation (6) reveals an important property of sampled data systems which is that the response characteristic for sampled sinusoidal inputs is periodic in the frequency domain with period T^{-1}.

The above derivation of Equation (6) is entirely analogous to the approach used with continuous systems where the characteristic defines the response to superposed complex frequencies. Here, the characteristic defines the response to sampled complex frequencies. It is worth noting that Equation (6) could have been deduced by substituting the analogue transfer function for a delay, that is, $e^{-j\omega T}$, wherever z^{-1} occurs.

2.3. Poles and zeros

As explained in the previous section, the system transfer function $H(z)$ can be expressed as the ratio of nth to mth order polynomials, in the variable z^{-1}, with real coefficients. Each polynomial may be factorised into n and m factors respectively and the values of z which make each numerator or denominator factor zero correspond to the zeros or poles, respectively, of the transfer function. Thus a purely recursive filter has only poles whilst a non-recursive filter has only zeros. If the positions of the poles and zeros are plotted in the complex z plane the resulting pattern is often very helpful in indicating the behaviour of the filter. Moreover, the pole-zero positions describe the transfer function completely except for a scaling factor.

Now it is well known, in the design of continuous systems, that a stable transfer function may not have any poles in the right-hand half of the p plane. As the imaginary axis of the p plane transforms into the unit circle of the z plane given by

$$|z| = 1$$

it is clear that a stable recursive filter may have poles only within the unit circle. Non-recursive filters, having no poles, are always stable.

It follows from the factorisation of $H(z)$ that the modulus and angle of $H(z)$ for any value of z are given by

$$|H(z)| = \frac{|z - z_{q1}| \text{X} |z - z_{q2}| \text{X} \dots}{|z - z_{p1}| \text{X} |z - z_{p2}| \text{X} \dots} = \frac{Q_1 Q_2 \dots}{P_1 P_2 \dots}$$

$$\phi = \alpha_1 + \alpha_2 + \dots - (\beta_1 + \beta_2 + \dots)$$

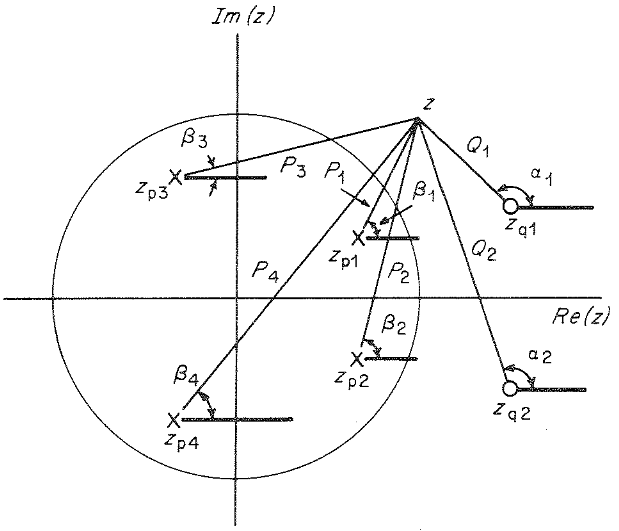

Fig. 2 - -Geometrical interpretation of $H(z)$

where z_{q1}, z_{q2} etc. are the zero positions and z_{p1}, z_{p2} etc. are the pole positions. This may be interpreted geometrically as in Fig. 2.

From this interpretation it should be clear that for example a low-pass behaviour will be obtained from a filter whose poles cluster near the sinusoidal frequency origin, i.e. the point (1,0) just as with continuous filters. However, there is a danger in pressing the analogies between continuous and discrete systems too far because although any point in the p plane transforms into one point in the z plane, Equation (5) shows that one point in the z plane transforms into an infinity of points in the p plane whose imaginary parts are separated by $2\pi T^{-1}$. Thus a digital filter, having poles and zeros in the z plane, which is based on those of an apparently equivalent continuous filter (through Equation (5)), will not have the same response characteristic as the continuous filter because of the presence of extra groups of poles and zeros in the p plane. In practice the responses for frequencies much less than the sampling frequency will be very similar if the poles and zeros are clustered very near the frequency origin. However, at frequencies near half the sampling frequency, the response of the digital filter will always depart from that of the continuous filter because at such frequencies the influence of the neighbouring pole-zero pattern (in the p plane) is as strong as that of the pattern near the origin.

2.4. Elemental forms of filter

Although the general transfer function $H(z)$ may have many poles and zeros, and therefore factors, it may be broken up into smaller units by factorisation or partial-fraction expansion to yield forms which can be cascaded or paralleled respectively. As the polynomial coefficients are real the poles and zeros must occur on either the real axis of z or in complex conjugate pairs. Thus the simplest units are those which are described by a single real pole or zero, or one conjugate pair. In the study of recursive filters the single pole or linear form and the double pole or quadratic

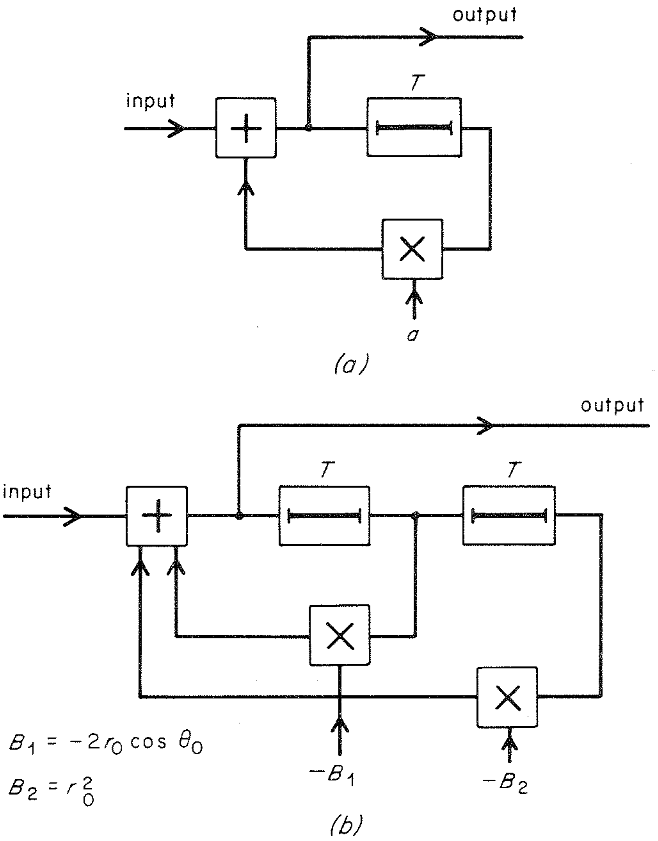

Fig. 3 - Circuits of first- and second-order recursive filters
(*a*) first-order (*b*) second-order

form of transfer function assume special significance and it is common, though by no means essential, to regard these forms as the basic 'building blocks' for more complex filters. The circuit arrangements of these filters are shown in Fig. 3. The corresponding transfer functions are

$$H_1(z) = 1/(1 - az^{-1})$$

and $$H_2(z) = 1/(1 - 2r_0 \cos\theta_0 z^{-1} + r_0^2 z^{-2})$$

from which it can be seen that the first-order pole is at $z = a$ and the second order poles are at $z = r_0 \exp \pm j\theta_0$. If the filters are stable then, of course, a and r_0 must be less than unity. These pole positions indicate that a first-order recursive filter has a low-pass behaviour whereas a second-order filter (if $4B_2 > B_1^2$) simulates a bandpass or single resonant circuit behaviour. This Report is mainly concerned with the first type of filter.

2.5. Quantisation effects

So far the discussion has assumed that the input and output samples and the coefficients A_i and B_i are all specified with complete accuracy. In practice, where digital arithmetic is used, all these quantities are specified with only limited accuracy and consequently the values of the computed samples differ from those predicted theoretically in the absence of quantisation.

The effects of finite accuracy can be divided into three categories. These are due to the quantisation of (1) the input samples, (2) the constants, and (3) the computed samples. The category 3 includes any intermediate values which form part of the final output value.

The effect of the filter on quantised input-samples can be analysed by assuming that the quantisation adds noise to the input signal with an r.m.s. value equal to $E_0/2\sqrt{3}$ where E_0 is the voltage corresponding to one quantum difference. As the filter is assumed linear the behaviour of the noise can be separated from that of the signal by superposition and thus the noise performance can be deduced assuming a pure noise input. In this way it can be shown that the steady-state mean square noise output of the filter is given by

$$\overline{e_o^2} = \overline{e_i^2} \sum_{k=0}^{\infty} h^2(kT)$$

where $\overline{e_i^2}$ is the m.s. input noise and $h(kT)$ is the filter impulse response.

For example the impulse response of the single pole filter of Fig. 3(*a*) is given by

$$h(kT) = a^k$$

so that $$\overline{e_o^2}/\overline{e_i^2} = \sum_{k=0}^{\infty} a^{2k} = 1/(1 - a^2)$$

If the pole is close to the unit circle, so that $a = 1 - \epsilon$ where ϵ is small, the ratio of the noise powers becomes $1/2\epsilon$ showing that the mean-square output noise is inversely proportional to the distance of the pole from the unit circle. However, the voltage gain at zero frequency is also inversely proportional to ϵ so that the ratio of signal-power to noise-power at low frequencies is also inversely proportional to ϵ. Thus the ratio increases as ϵ decreases. This is, after all, quite reasonable because as ϵ decreases the filter becomes more selective.

The quantisation of the constant parameters of the filter, A_i and B_i, merely shifts the pole and zero positions from their intended positions to new ones. In non-recursive designs the effect of this movement is not serious and the resulting deviation from the designed specification can be easily calculated. In recursive designs, however, the quantisation may move a pole from inside the unit circle to the boundary, so making the filter unstable.

For example the value of a in the single pole filter of Fig. 3 needs to be specified with increasing accuracy as it approaches unity. For the second order filter in Fig. 3 it can be shown[3] that

$$\Delta r_o = \Delta B_2/2r_o, \quad \Delta\theta_o = \frac{\Delta B_2}{2r_o^2 \tan\theta_o} + \frac{\Delta B_1}{2r_o \sin\theta_o}$$

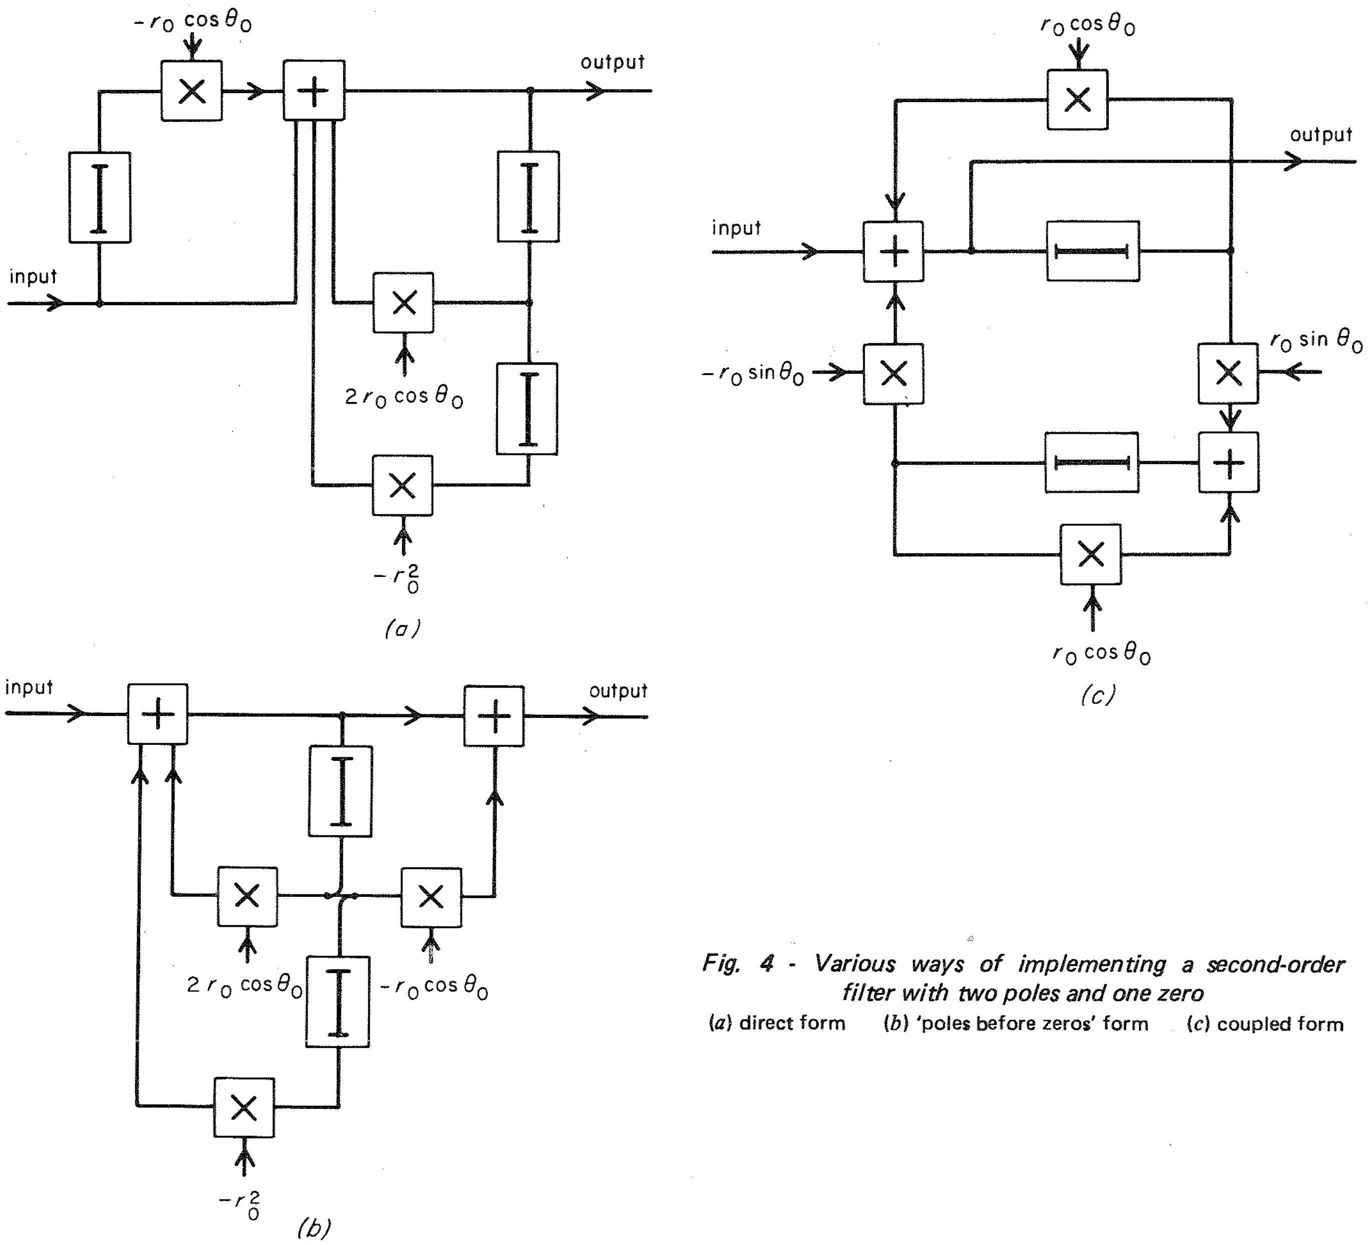

Fig. 4 - Various ways of implementing a second-order filter with two poles and one zero
(*a*) direct form (*b*) 'poles before zeros' form (*c*) coupled form

so that if r_0 is close to unity and therefore also B_2, this coefficient needs to be specified with increasing accuracy as r_0 approaches unity. Further, for given inaccuracies in B_1 and B_2 the error in θ_0 is worse for low values of θ_0. Filters requiring low values of θ_0 arise when the sampling rate is made very high so that the poles cluster near the frequency origin. In such cases, therefore, the filter coefficients must be very accurate.

The effects caused by quantising the results of additions and multiplications are much more difficult to analyse because in recursive designs the errors so introduced are perpetuated by the action of the filter. However, it is possible to analyse the effects by setting up a noise model wherein an equivalent noise source is associated with each addition or multiplication (only the latter need be considered in filters using fixed-point arithmetic). With the aid of such a model it can be shown that the many differing forms in which a transfer function can be realised perform very differently when quantisation is taken into account.

For example, Fig. 4 shows various ways of implementing a transfer function consisting of two poles at $r_0 \exp \pm j\theta_0$ and a zero at $r_0 \cos\theta_0$. Arrangements (*b*) and (*c*) are said to be canonic because they contain as many delay elements as the order of the transfer function (that is, the number of poles). Arrangement (*a*) as in Fig. 1 is, for obvious reasons, called the direct form. In fact, arrangement (*c*) which is called the coupled form has the best noise performance.

As a general rule the direct form never gives the best noise performance. Moreover, as the order of the filter increases, the positions of the poles and zeros become more and more sensitive to slight changes in the coefficient values of the direct form. Therefore it is always preferable to implement high-order filters through serial or parallel combinations of low-order filters.

A further degree of freedom can be introduced into the noise analysis by allowing intermediate results to be stored with greater accuracy than the final answer. This,

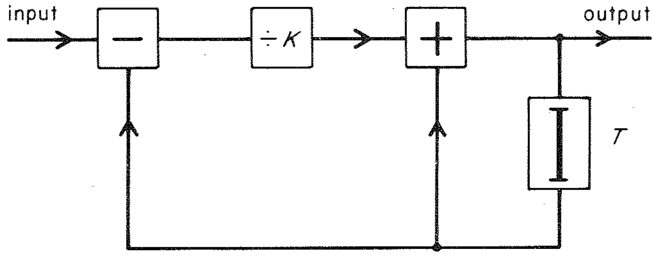

Fig. 5 - Modified form of first-order recursive filter

of course, means greater register lengths and longer computation times, and the whole art of digital-filter design, once the pole and zero positions have been specified, is to realise a transfer function which has the best noise performance consistent with the minimum requirements of register length.

3. The first-order recursive filter — theoretical

3.1. The pole plot

The filter chosen for study was the simplest form of recursive filter shown in Fig. 3(*a*), but modified as in Fig. 5. In its modified form the filter always has unity gain at zero frequency which avoids the embarrassment of very high gains as a tends to unity. If the pole positions of the filters in Figs. 3(*a*) and 5 are identical it is easily shown that

$$a = 1 - K^{-1} \text{ or } K = (1 - a)^{-1}$$

The corresponding transfer function for Fig. 5 is

$$H(z) = 1/[K - (K-1)z^{-1}] \tag{7}$$

which shows the unity-gain behaviour at zero frequency, for which $z = 1$.

3.2. The frequency characteristic

The response characteristic for sampled sinusoidal inputs is given by $H(e^{j\omega T})$ and as Equation (7) shows that it is more convenient to work with the reciprocal characteristic $H^{-1}(e^{j\omega T})$, the result is

$$H^{-1}(e^{j\omega T}) = K - (K-1)e^{-j\theta} \tag{8}$$

where $\theta = \omega T$. θ is in fact the phase shift through the delay element.

Equation (8) can be interpreted geometrically as in Fig. 6 in which frequency is linearly scaled round the circumference of the circle of radius $K - 1$. From Equation (8) the modulus and phase of $H^{-1}(e^{j\omega T})$ are given by

$$|H^{-1}|^2 = 1 + 2K(K-1)(1 - \cos\theta)$$

$$-\phi = \tan^{-1}\left\{\sin\theta/(1 + [K-1]^{-1} - \cos\theta)\right\}, K \neq 1$$

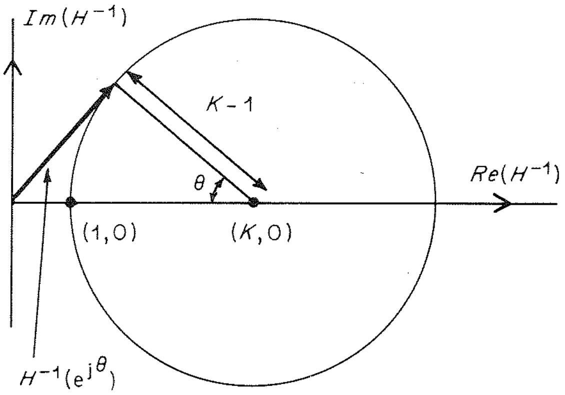

Fig. 6 - Geometrical interpretation of reciprocal transfer function of filter in Fig. 5.

These relationships are plotted in Fig. 7 for various values of K where it will be seen that the functions are periodic as expected. As the value of K is increased the amplitude/frequency characteristic becomes more and more peaked at zero frequency showing that the filter becomes more selective. Note that to obtain an equivalent sharpness of response using a non-recursive filter many more delay elements and associated multipliers would be required. The 3 dB point of the characteristic is given by

$$2K(K-1)(1 - \cos\theta) = 1$$

For $\theta < 0{\cdot}5$ a good approximation to $|H^{-1}|^2$ is given by

$$|H^{-1}|^2 \doteqdot 1 + K(K-1)\theta^2 \tag{9}$$

Consequently the 3 dB point is given by

$$\theta = \omega T \doteqdot [K(K-1)]^{-\frac{1}{2}}$$

$$\rightarrow K^{-1} \text{ for large } K.$$

3.3. The pulse response

The pulse response of the filter can be written down by inspection noting that the output consists of the present input scaled by K^{-1} and the previous output scaled by $1 - K^{-1}$. Thus for the input sequence 1, 0, 0, ... the output is K^{-1}, $K^{-1}(1 - K^{-1})$, $K^{-1}(1 - K^{-1})^2$, . . . which will be recognised as a sequence of samples having an exponential decay. The filter therefore behaves like a simple RC integrator in the sense that the output values are samples of the corresponding analogue pulse-response. This similarity is termed impulse invariance and is sometimes the basis on which digital filters are designed.

3.4. The relationship with the analogue integrator

The reciprocal transfer function of a continuous RC integrator is

$$H^{-1} = 1 + j\omega\tau$$

whence

$$|H^{-1}|^2 = 1 + \omega^2\tau^2 \tag{10}$$

Comparing Equations (9) and (10) the equivalent analogue integrator time constant, τ, is given by

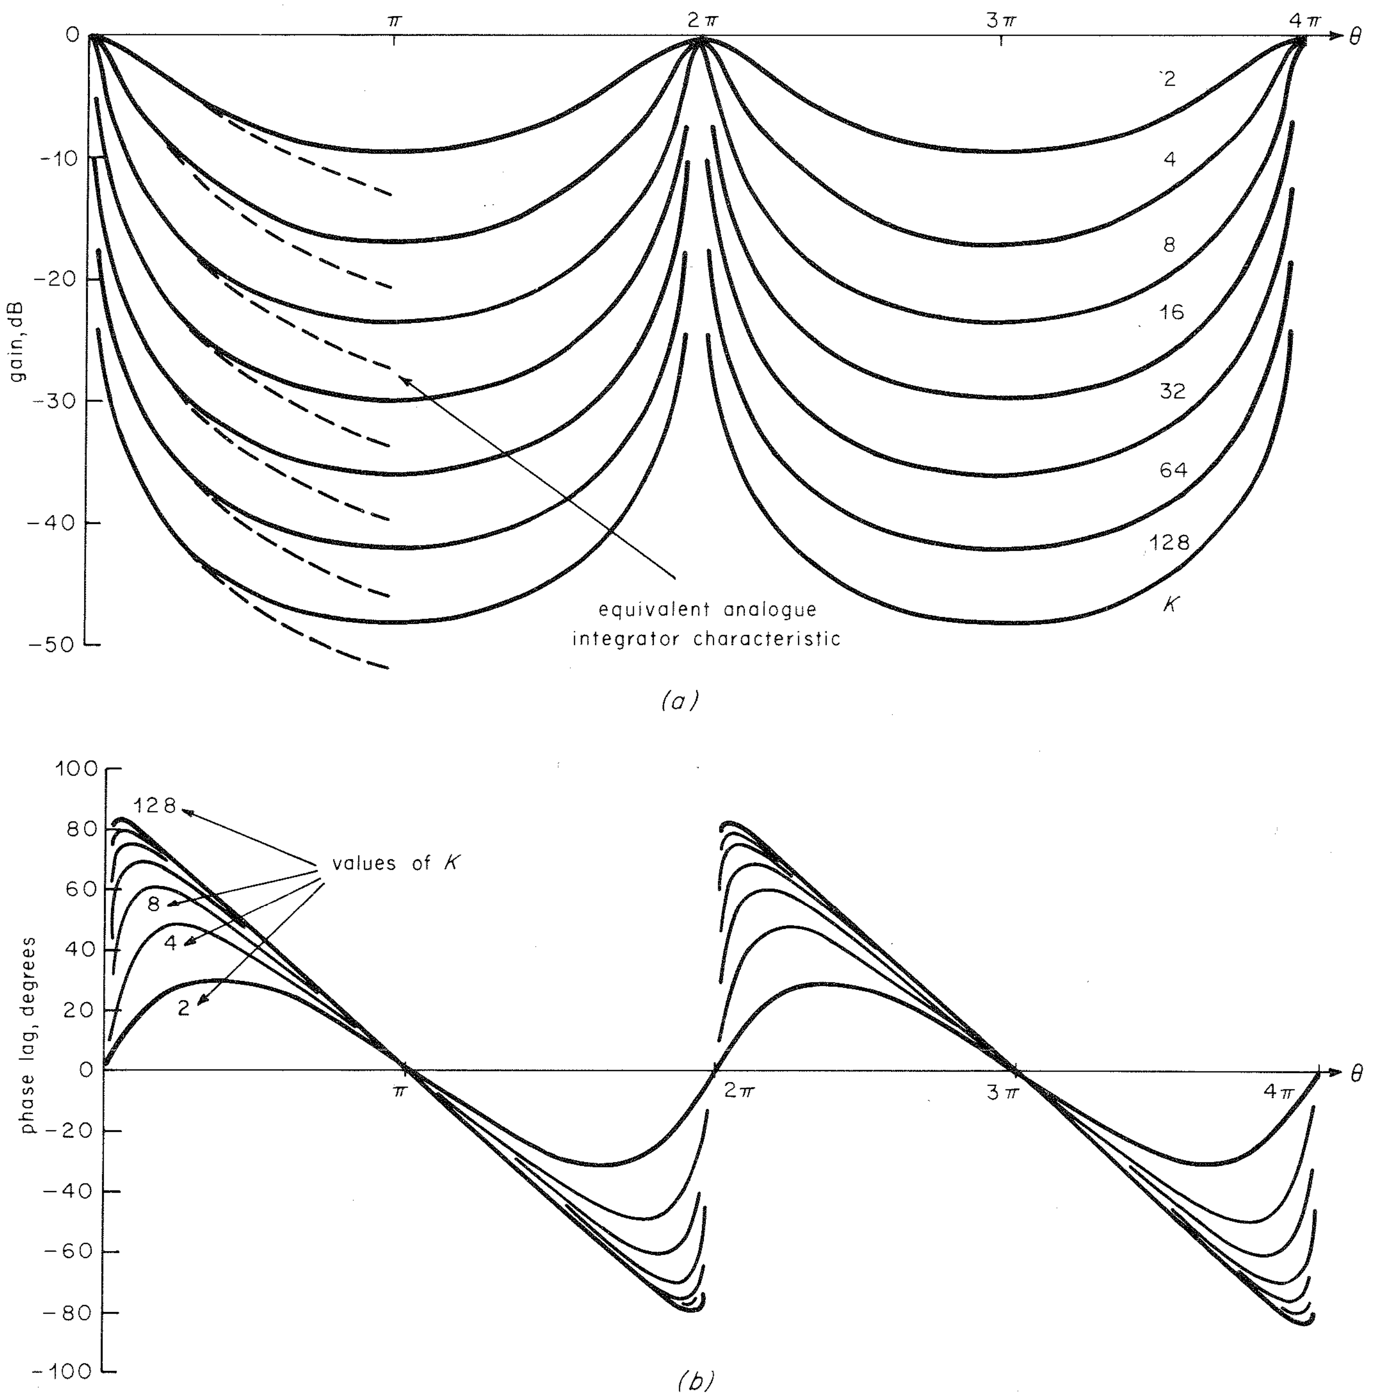

Fig. 7 - Characteristics of filter in Fig. 5
(*a*) amplitude/frequency (*b*) phase/frequency

$$\tau = [K(K-1)]^{\frac{1}{2}}T \qquad (11)$$

Thus as K tends to infinity the equivalent time constant tends to K times the sampling period.

The form of Equation (9) also shows that the responses of the digital and analogue circuits are substantially equal for $\theta < 0{\cdot}5$ as can be seen in Fig. 7. But as ω tends to infinity the analogue integrator attenuation tends to infinity while the digital filter offers only a finite maximum attenuation of odd multiples of half the sampling frequency. This maximum is $20\log_{10}(2K-1)$ whereas the attenuation of the analogue integrator at the same frequency is $10\log_{10}[1+K(K-1)\pi^2]$ using Equation (11). Thus for large values of K the attenuation of the analogue circuit exceeds that of the digital filter at half the sampling frequency by

$$20\log_{10}K\pi - 20\log_{10}2K \doteqdot 4\text{ dB.}$$

3.5. The effects of quantisation

Only the effect of quantising computed values will be considered in detail here. For instrumental reasons it is convenient to quantise values of K to integer powers of 2.

The effects of quantisation can be illustrated by first assuming that computed values at each point in the filter are specified with the same number of bits. In a fixed-point filter multiplication by a constant produces, in general, further bits of lower significance. These bits can be discarded by rounding or truncation. The effect of rounding is always to increase or decrease a number to the nearest appropriate level but the result of truncation is linked with the way in which negative numbers are represented. In offset binary representation the effect of truncation is to decrease a positive number but increase the modulus of a negative number to the appropriate level of significance.

In the filter of Fig. 5 the division by K produces extra bits and so the number which enters the adder can be rounded or truncated. The result of this operation is best appreciated by considering the pulse response of the filter. From Section 3.3 the response in the absence of quantisation, that is with an infinite number of bits, would be an exponential decay to zero with a time constant given by Equation (11).

First consider rounding. The signal at any point in the filter can be expressed by an integer. Suppose at any time after the impulse has occurred the output of the filter is b. Then to form the next output, the number b is subtracted from the input (zero) to produce $-b$, divided by K to produce $-b/K$, rounded to $[-b/K]$ * and finally added to b to produce $b + [-b/K]$. Thus the output decreases by $[b/K]$. Now this sequence of events will continue until b reaches $K/2$. If the rounding is such that a residue of ½ is rounded up then b/K becomes 1 and the output decreases once more. Thereafter $[b/K]$ is zero and the output does not decrease further. For example if K is 8 and b is initially 7 the output sequence obtained is

$$7, 6, 5, 4, 3, 3, 3, \ldots$$

and thus there is a constant error. This constant error is characteristic of first order filters and is referred to as a 'deadband' error. Higher order filters have more complicated behaviour in their deadbands.

The response to a negative impulse would end with an identical error of opposite sign because of the symmetrical effect of rounding on positive and negative numbers. Moreover by similar reasoning it can be shown that the step response ends with the same error such that a positive step ends with a negative error and a negative step ends with a positive error.

Consider now truncation. As before an output of b is followed by an output of $b + \{-b/K\}$.** Now if b is positive and less than $K/2$ the quantity $-b/K$ will always be truncated to -1 however small the value of b. Thus the output of the filter continues to decrease by one after b passes $K/2$. When the output reaches zero there is an abrupt termination of the sequence with no error. If, however, b is negative and numerically less than K the quantity $-b/K$ is truncated to zero, and the filter output cannot increase further. For example if K is 8 and b is initially -9 the output sequence is

$$-9, -8, -7, -7, -7, \ldots$$

Thus with truncation the response to a positive impulse ends with no error but the response to a negative impulse ends with an error of approximately twice that due to rounding.

By similar reasoning it is possible to show that the response to a negative step ends with no error but the response to a positive step ends with the same error as above. Thus the asymmetrical behaviour of truncation is reflected in the responses.

However, these steady-state errors can be reduced or eliminated by storing the computed results more accurately. Suppose therefore the capacities of the store, adder and subtractor are increased by q extra bits of lower significance. The discussion can be framed in exactly the same terms as before except that b now represents a number which is 2^q times smaller. As a result the errors are also 2^q times smaller.

Mathematically the error due to rounding, E_R, without the extra bits can be expressed as

$$|E_R| = 2^{p-1} - 1, \quad p>0$$

where $K = 2^p$.

The error due to truncation for a negative pulse, E_{T-}, is given by

$$-E_{T-} = 2^p - 1$$

where the negative suffix denotes a negative-going pulse. With the extra bits the errors become

$$|E_R| = 2^{-q}(2^{p-1} - 1) \tag{12a}$$

$$-E_{T-} = 2^{-q}(2^p - 1) \tag{12b}$$

These relationships are tabulated in Table 1.

If $|E_R|$ is less than half an input quantum, Table 1 shows that q must equal p unless p is 1 in which case q can be zero. If $-E_{T-}$ satisfies the same criterion Table 1 shows that q must equal $p + 1$. Thus if p is large we obtain the important result that division by K in a first-order filter of the form of Fig. 5 requires an expansion of the arithmetic capacity by approximately $\log_2 K$ bits.

There is, however, a further complication introduced by including the extra bits. If the filter does not form part of a more complex arrangement the accuracy required at the output will only be the same as that at the input. (It should be recalled that because the gain is unity at zero frequency the output lies in the same range as the input.) So the extra bits in the signal at the input of the delay must

---

* where [] denotes rounding.

** where {} denotes truncation.

TABLE 1

Errors (in Terms of Input Quanta) in Filter Output Assuming no Output Truncation

E_R								$-E_{T-}$
p \ q	0	1	2	3	4	5	6	q / p
	Values of E_R and $-E_{T-}$ ($E_{T+} = 0$)							
1	0	0	0	0	0	0	0	0
2	1	$^1/_2$	$^1/_4$	$^1/_8$	$^1/_{16}$	$^1/_{32}$	$^1/_{64}$	1
3	3	$^3/_2$	$^3/_4$	$^3/_8$	$^3/_{16}$	$^3/_{32}$	$^3/_{64}$	2
4	7	$^7/_2$	$^7/_4$	$^7/_8$	$^7/_{16}$	$^7/_{32}$	$^7/_{64}$	3
5	15	$^{15}/_2$	$^{15}/_4$	$^{15}/_8$	$^{15}/_{16}$	$^{15}/_{32}$	$^{15}/_{64}$	4
6	31	$^{31}/_2$	$^{31}/_4$	$^{31}/_8$	$^{31}/_{16}$	$^{31}/_{32}$	$^{31}/_{64}$	5
7	63	$^{63}/_2$	$^{63}/_4$	$^{63}/_8$	$^{63}/_{16}$	$^{63}/_{32}$	$^{63}/_{64}$	6

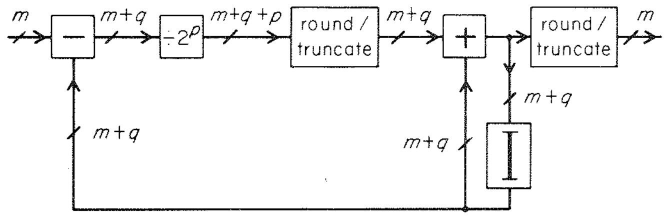

Fig. 8 - Filter model showing numbers of bits at each stage

be discarded when forming the output. Here again rounding or truncation can be applied. This is illustrated in Fig. 8 for the filter of Fig. 5.

Because of this further stage of quantisation it is possible to define four situations. These arise from a rounding or a truncation within the loop combined with a rounding or a truncation outside it. The corresponding errors can be designated E_{RR}, E_{RT}, E_{TR} and E_{TT} and where truncation is involved these errors will be different for positive and negative impulses. The complete behaviour of these errors is summarised in Table 2. It must be remembered that the error for a positive step is the same as that for a negative impulse and vice versa.

From Table 2 it can be seen that E_{RT} and E_{TT} can never equal zero for both polarities of signal (except for $p = 1$). E_{RR} and E_{TR} can, but in such a case the E_{TR} situation requires $p + 1$ extra bits whereas the E_{RR} situation requires only p bits. Moreover if less than p bits are added only E_{RR} is symmetrical for both polarities of signal. Thus rounding inside and outside the loop gives the best results.

In summary, the characteristic deadbands of a first order filter which correspond to steady-state errors can be eliminated by expanding the arithmetic capacities of the filter elements. The errors corresponding to various degrees of expansion have been tabulated and are seen to depend on whether rounding or truncation is applied to the numbers. Rounding gives the better result but is instrumentally more complicated than truncation which is very simple.

4. The design of a first-order recursive filter

4.1. The sampling frequency

A filter based on the aforementioned principles was designed to operate on video signals. The Nyquist rate for such signals, being twice the highest component frequency, is 11 MHz, but a sampling frequency which is convenient for some applications is thrice colour subcarrier frequency, namely 13·300 . . . MHz, which is well above the Nyquist limit. The filter was designed to work at this latter sampling frequency. The corresponding sampling interval is approximately 75 ns so that if K lies in the range 1 – 100 the maximum time constant obtainable would be, from Equation (11), of the order of 7·5 µs. The corresponding minimum 3 dB frequency would be about 21 kHz because, for the chosen sampling frequency, the 3 dB frequency is given approximately by

$$f_{3\text{ dB}} = 2{\cdot}1/[K(K-1)]^{1/2}\ \text{MHz}$$

$$\doteqdot 2{\cdot}1/K\ \text{MHz for } K \text{ large.}$$

As the frequency response characteristic is periodic having a period equal to the sampling frequency, the frequency of maximum attenuation is half the sampling frequency i.e. at approximately 6·65 MHz.

4.2. The number of bits

The filter was designed to accept and deliver a stan-

TABLE 2

Errors (in Input Quanta) in Filter Output when Output has Same Number of Bits as Input

$|E_{\mathrm{RR}}|$

q p	0	1	2	3	4	5	6
1	0	0	0	0	0	0	0
2	1	1	0	0	0	0	0
3	3	2	1	0	0	0	0
4	7	4	2	1	0	0	0
5	15	8	4	2	1	0	0
6	31	16	8	4	2	1	0
7	63	32	16	8	4	2	1

$E_{\mathrm{RT+}}$

q p	0	1	2	3	4	5	6
1	0	0	0	0	0	0	0
2	1	0	0	0	0	0	0
3	3	1	0	0	0	0	0
4	7	3	1	0	0	0	0
5	15	7	3	1	0	0	0
6	31	15	7	3	1	0	0
7	63	31	15	7	3	1	0

$E_{\mathrm{RT-}}$

q p	0	1	2	3	4	5	6
1	0	0	0	0	0	0	0
2	−1	−1	−1	−1	−1	−1	−1
3	−3	−2	−1	−1	−1	−1	−1
4	−7	−4	−2	−1	−1	−1	−1
5	−15	−8	−4	−2	−1	−1	−1
6	−31	−16	−8	−4	−2	−1	−1
7	−63	−32	−16	−8	−4	−2	−1

$E_{\mathrm{TR-}}$ $(E_{\mathrm{TR+}} = 0)$

q p	0	1	2	3	4	5	6
1	−1	−1	0	0	0	0	0
2	−3	−2	−1	0	0	0	0
3	−7	−4	−2	−1	0	0	0
4	−15	−8	−4	−2	−1	0	0
5	−31	−16	−8	−4	−2	−1	0
6	−63	−32	−16	−8	−4	−2	−1
7	−127	−64	−32	−16	−8	−4	−2

$E_{\mathrm{TT-}}$ $(E_{\mathrm{TT+}} = 0)$

q p	0	1	2	3	4	5	6
1	−1	−1	−1	−1	−1	−1	−1
2	−3	−2	−1	−1	−1	−1	−1
3	−7	−4	−2	−1	−1	−1	−1
4	−15	−8	−4	−2	−1	−1	−1
5	−31	−16	−8	−4	−2	−1	−1
6	−63	−32	−16	−8	−4	−2	−1
7	−127	−64	−32	−16	−8	−4	−2

dard PCM video signal of 8-bit word length. It was thought worthwhile to experiment with values of K up to 128, giving a maximum attenuation of 48 dB. This value of K is convenient for instrumental reasons because if K is restricted to integer powers of 2, division by K is simply carried out by shifting data to lower levels of significance; a shift of up to 7 places can be provided by an 8 way data selector.

The preceding discussion of quantisation effects indicate that, for such a range of values for K, a word length of 15 bits would be needed in the store, adder and subtractor to avoid all traces of deadband effects. In addition it was thought desirable to make provision for overflow when using the filter as part of a more complex arrangement to be described later. Thus a word length of 16 bits was chosen which was also instrumentally convenient. Of these bits the two most significant were designated as overflow bits, the next eight as signal-range bits and the last six as 'accuracy' bits.

4.3. The need for ECL

The need to carry out 16-bit addition, subtraction, variable shift and storage, all within 75 ns, precluded the use of any form of TT logic. The fastest family of TTL can do these operations in a typical time of 68 ns but in a 'worst case' time of 107 ns, even assuming full 'look-ahead' carry generation is used. Moreover it was planned to expand the variable shift operation to include an alternative addition/subtraction process. This modification is described in a companion Report.[4]

As a result of this demand for speed, only emitter-coupled logic could be used for the aforementioned processes. As the PCM signal is normally generated and decoded in TTL form, logic-level translators were included preceding and following the filter; however, their speed of operation was of no account because they lay outside the loop. The emitter-coupled logic performed the required operations in a 'worst-case' time of 48 ns, giving a generous margin for subsequent modifications.

4.4. Circuit details

The construction of the adder, subtractor and store are quite straightforward and need no further elaboration. The variable data shifter comprised 17 8-way data selectors connected to the output of the subtractor in such a way that successive selector inputs were connected to subtractor outputs of successively higher significance. In both TTL and ECL logic families the result of a subtraction is expressed in offset binary, that is, the extra most significant bit generated is 'one' for positive numbers and 'zero' for negative numbers. This means that the number zero lies in the middle of the range and is therefore 'offset'. Thus a small positive number begins with 100 . . . . whereas a small negative number begins with 011 . . . . Division by 2^p, resulting in a shift of p places to the right, requires that the p places to the right of the first, i.e. sign digit be filled with the complement of the sign digit. This quantity must therefore be fed to all remaining inputs of the data selectors.

Once the sign digit has been used in this way it can be discarded. Fig. 9 shows why this is so. Two sequences are shown which represent the way in which numbers are expressed at different points in the filter. At any instant the present input and previous output sample values are x and y respectively. At (a) x is assumed to be greater than y, at (b) x is assumed to be less than y. In either case the final result must lie between x and y because it is a linear weighting of x and y with coefficients K^{-1} and $1 - K^{-1}$.

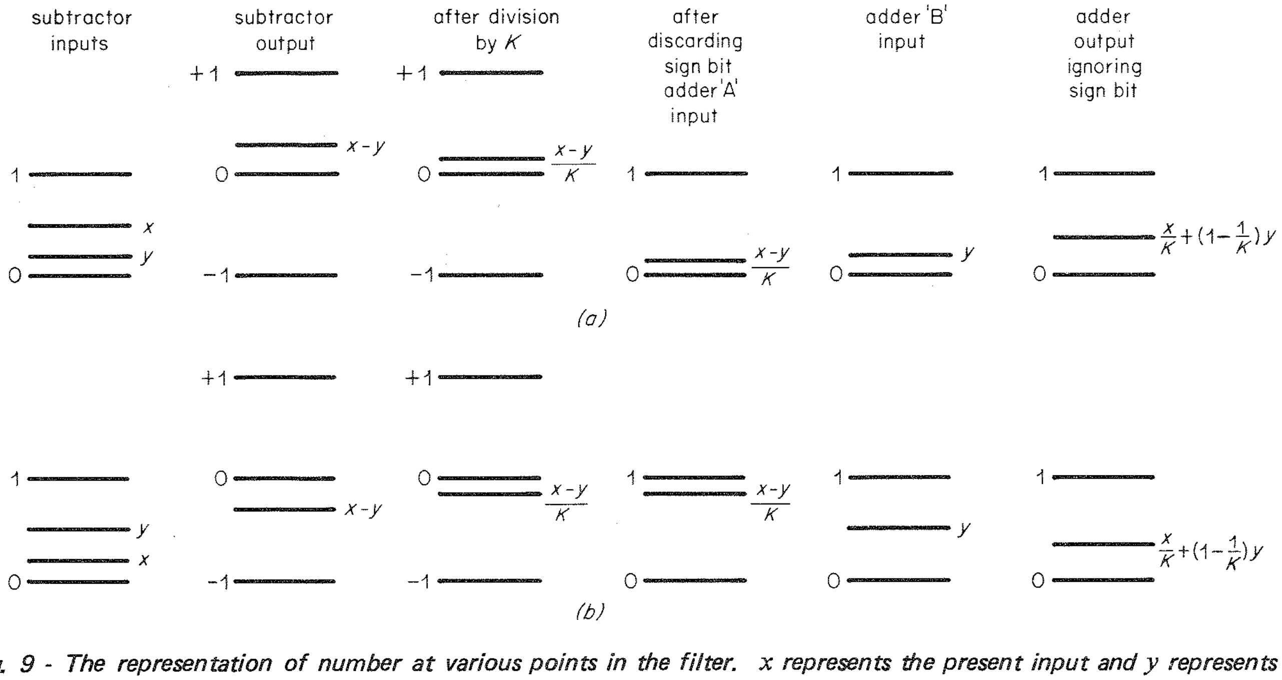

Fig. 9 - The representation of number at various points in the filter. x represents the present input and y represents the previous output

(a) $x > y$ (b) $x < y$

Rounding is performed by adding in, at the level of the least significant bit, the most significant of the bits discarded by shifting. Thus if a number ending in 101 is shifted three places a 1 is added in; it can easily be verified that this works for both positive and negative numbers in offset binary notation. However, a residue of exactly ½ is rounded up. The addition is conveniently performed in the adder following the data shifter. Thus the shifter must carry forward an extra 17th bit which is fed into the carry input of the following adder.

Truncation is used to obtain a standard 8-bit signal from the loop. This is because it is easier to instrument than rounding which would require a further adder. As a result the final value can be in error by one eighth-bit (see the last column of $E_{\mathrm{RT}+}$ and $E_{\mathrm{RT}-}$ in Table 2).

4.5. Observed performance

The amplitude/frequency response characteristic was as predicted, i.e. as shown in Fig. 7. However, at high frequencies the output signal was so small that it only explored a few quantum levels and its amplitude was therefore difficult to measure. Fig. 10 shows the effect of discarding bits from the circulating loop. For this test the input was set at half the input range, peak-to-peak (neglecting overflow), and K was set to 128; this combination effectively demonstrates the need for the extra bits when the filtering is severe. The pulse response was also as predicted, i.e. a decaying exponential with the appropriate time constant. Fig. 11 shows the effect of discarding bits from the circulating loop. This introduces a final error, as can be seen.

5. The vertical first-order recursive filter — theoretical

5.1. The relationship with the horizontal filter

The term vertical is used to distinguish the filter from the preceding case. Consider a two-dimensional array of samples $\{x_{kl}\}$. A horizontal filter operates upon samples of a row, i.e. the samples x_{kl}, $x_{k,l+1}$, $x_{k,l+2}$ etc. Correspondingly, the action of a vertical filter relates to the samples of a column i.e. the samples $x_{k,l}$, $x_{k+1,l}$, $x_{k+2,l}$ etc. This behaviour can be dealt with using the two-dimensional z-transform. By extension this can be defined as

$$X(z_1, z_2) = \sum_{k=0}^{\infty}\sum_{l=0}^{\infty} x(kT_1, lT_2) z_1^{-k} z_2^{-l}$$

where the sample x_{kl} occurs at the point with Cartesian coordinates (lT_2, kT_1). In a way entirely analogous to one-dimensional analysis the output of a two-dimensional filter can be found by taking the two-dimensional z-transform of the input, multiplying by the z-transform of the filter and taking the inverse transform. Thus

$$Y(z_1, z_2) = H(z_1, z_2)X(z_1, z_2)$$

where the filter transfer function can be expressed as

$$H(z_1, z_2) = \sum_{i=0}^{n_1}\sum_{j=0}^{n_2} A_{ij} z_1^{-i} z_2^{-j} \Bigg/ \sum_{i=0}^{m_1}\sum_{j=0}^{m_2} B_{ij} z_1^{-i} z_2^{-j}$$

For a purely horizontal filter.

$$A_{ij} = B_{ij} = 0, \qquad i \neq 0$$

and the transfer function becomes

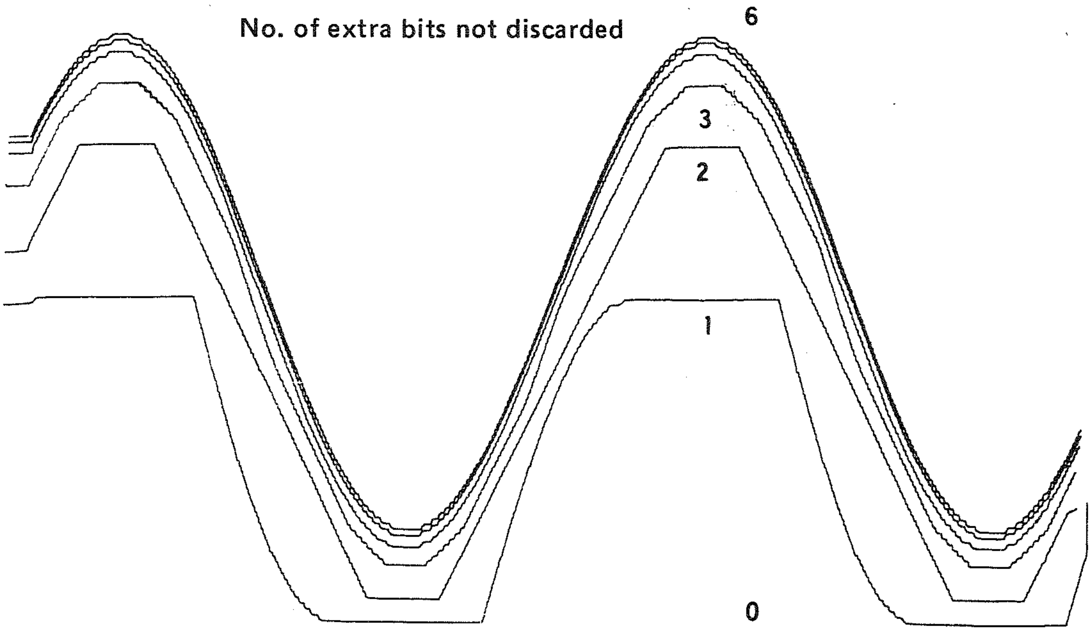

Fig. 10 - The effect on the sinewave response of discarding bits from the circulating loop

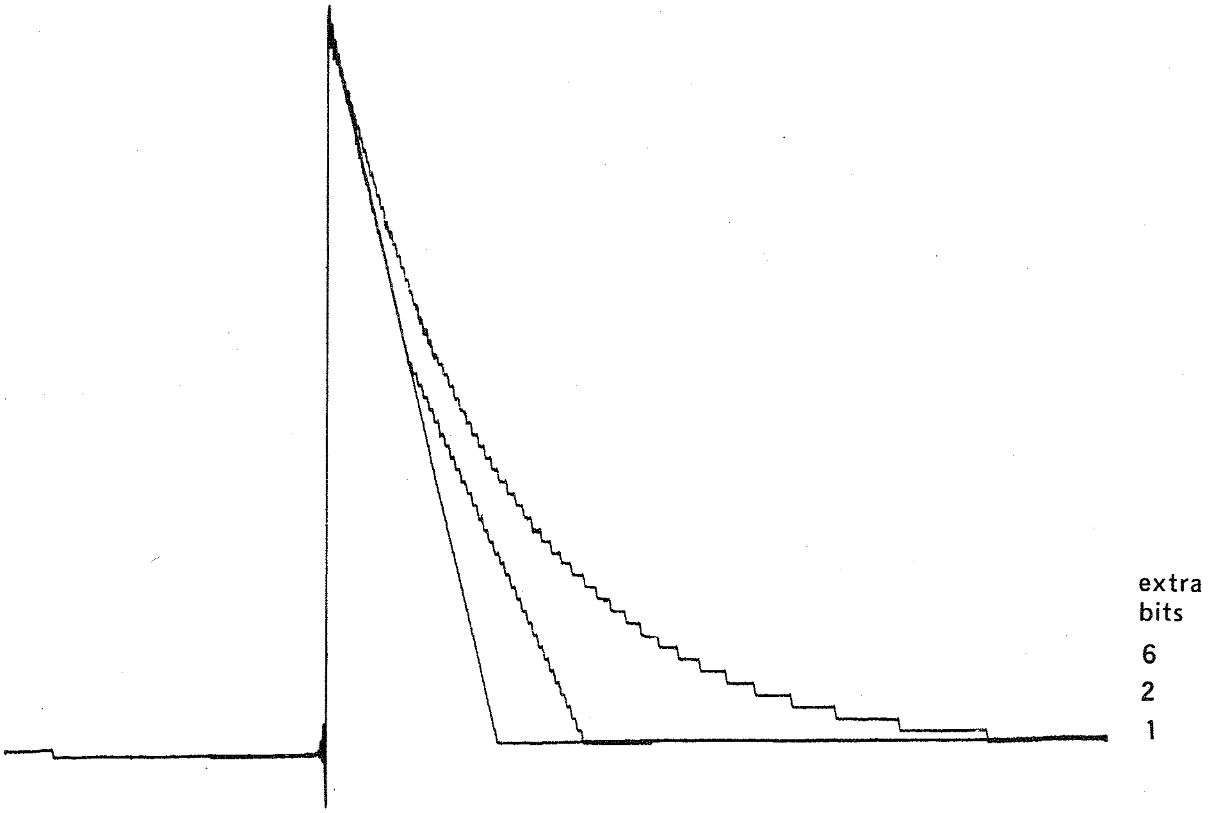

(*a*)

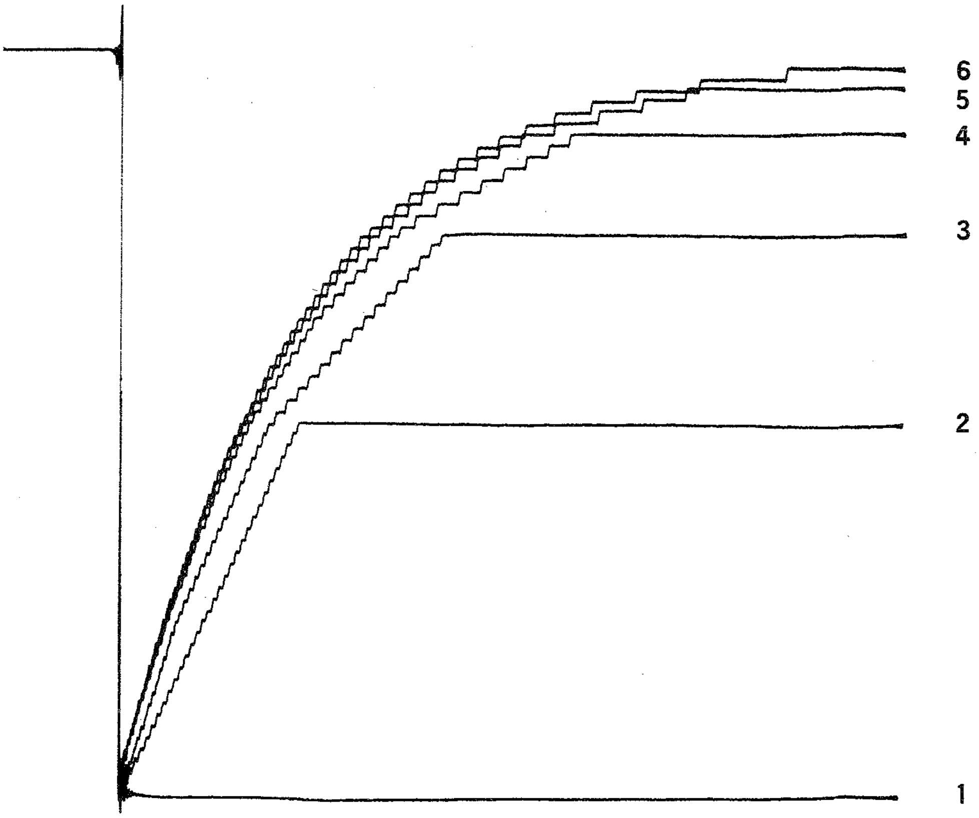

(*b*)

Fig. 11 - The effect on the pulse response of discarding bits from the circulating loop
(*a*) positive pulse (*b*) negative pulse

$$H(z_1, z_2) = \sum_{j=0}^{n_2} A_{0j} z_2^{-j} \bigg/ \sum_{j=0}^{m_2} B_{0j} z_2^{-j}$$

which should be compared with Equation (4). For a purely vertical filter

$$A_{ij} = B_{ij} = 0, \quad j \neq 0$$

and the transfer function becomes

$$H(z_1, z_2) = \sum_{i=0}^{n_1} A_{i0} z_1^{-i} \bigg/ \sum_{i=0}^{m_1} B_{i0} z_1^{-i}$$

It should therefore be clear that a purely vertical filter can be derived from the corresponding horizontal filter merely by interchanging the roles of i and j.

In a television system the two dimensional array of samples is converted by the scanning action into a one-dimensional array. If there are M samples per line we may write

$$x_{ij} = x_{j+Mi}$$

Thus a first-order vertical filter which relates $x_{i+1,j}$ to x_{ij} may also be considered, on a one-dimensional basis, to relate x_{j+M} to x_j. In other words a first-order vertical filter is a special form of Mth order horizontal filter with the $(M-1)$th, $(M-2)$th etc. terms absent.

5.2. The pole plot

Considered on a one-dimensional basis the transfer function of the first-order vertical recursive filter is given by

$$H(z) = 1/(1 - az^{-M})$$

This function has M poles given by the roots of the equation

$$z^M = a$$

i.e. $\quad z = a^{1/M}\,\mathrm{cis}(2r\pi/M),\; r = 0, 1, 2, \ldots .$

Thus the roots lie equi-spaced on a circle of radius $a^{1/M}$ as shown in Fig. 12.

The low-pass behaviour of the filter is not immediately obvious from the pole plot of Fig. 12. The repetitive nature of the pattern merely seems to indicate that the response characteristic has a fine structure with no particular emphasis on high or low frequencies.

Consider now the transfer function on a two-dimensional basis. This is given by

$$H(z_1, z_2) = 1/(1 - az_1^{-1})$$

with the corresponding pole plot as in Fig. 13. The low-pass behaviour is immediately clear.

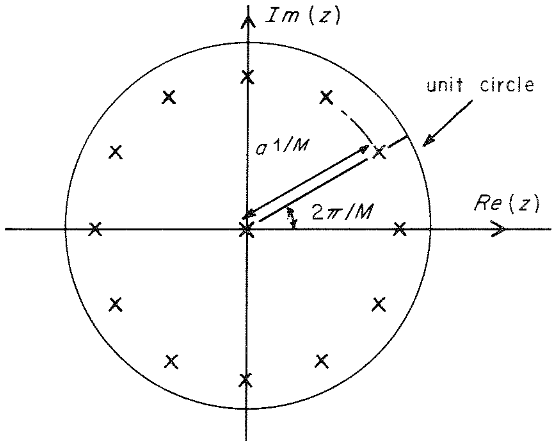

Fig. 12 - The pole plot of the vertical filter considered on a one-dimensional basis

If we add more terms to the vertical filter we create further poles in the z_1 plane and the resulting pattern appears as a repetitive fine structure in the one-dimensional plane of Fig. 12. Thus the plot of Fig. 13 contains all the information in Fig. 12 and is therefore a better representation.

5.3. The frequency characteristic

It follows from the above treatment that the response characteristic of the vertical filter is the same for purely vertical frequencies as that of the horizontal filter for purely horizontal frequencies i.e. that of Fig. 7. But θ now represents the phase difference across a vertical sampling element instead of a horizontal element. Moreover the response characteristic corresponding to any spatial frequency is that of Fig. 7 applied to the resolved vertical component. For example a purely horizontal frequency has no vertical component and therefore the filter attenuation is zero for such frequencies.

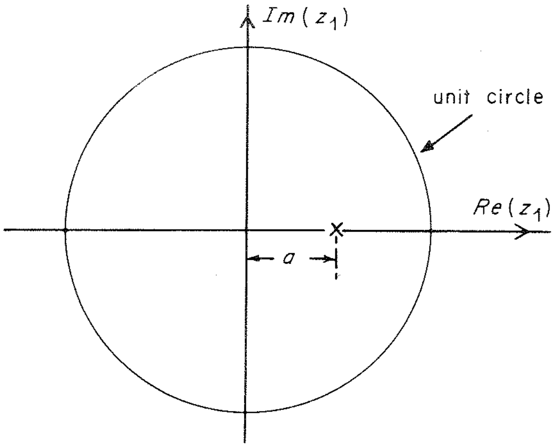

Fig. 13 - The pole plot of the vertical filter considered on a two-dimensional basis

At this point caution must be exercised. For an interlaced system the vertical sampling element corresponds to the pitch of the lines of a field. Thus the filter response characteristic for an input of 312½ c/p.h. will be the same as for zero frequency because, for this input, the lines of a field are identical. The vertical frequency for maximum attenuation is therefore 156¼ c/p.h.

This behaviour can also be seen from the temporal frequency response characteristic. This characteristic can be obtained by putting

$$\theta = 2\pi f T_L$$

in Equation (8) where T_L is now a line period. This gives a response characteristic which repeats at multiples of line frequency giving maximum attenuation at odd multiples of half line frequency. It is these latter frequencies which, in an interlaced system, correspond to resolved vertical components of 156¼ c/p.h.

A vertical filter having maximum attenuation at 312½ c/p.h., i.e. based on a sampling interval equal to the pitch of the picture lines, would involve a field delay and thus be considerably more complex.

5.4. The pulse response

The pulse response of the vertical filter, like the frequency response characteristic, can be derived from that of the corresponding horizontal filter. It is thus a vertically decaying exponential with a 'time' constant given by Equation (11) where T now represents the pitch of the lines of a field.

It follows from this behaviour that the action of the filter can be likened to that of an analogue vertical integrator; that is, the displayed result is the same as would have been achieved had the scanning been vertical and the signal passed through an integrator.

6. The design of a vertical first-order recursive filter

6.1. Circuit details

The preceding treatment showed the relationship between the horizontal and vertical filters. Implicit in this relationship was the assumption of an integral number of samples in each scanning line giving an orthogonal Cartesian array of sample points. This assumption makes the realisation of a purely vertical filter easy, although it is possible to realise such a filter based on a non-orthogonal array; however, the two-dimensional difference equations would then contain cross-terms and the filter would become considerably more complex.

In practical terms the simple vertical difference equation can be realised using a line store clocked with a line-locked frequency. The rest of the filter is as in Fig. 8 so that, for example, a 16-bit line store is needed. This can be formed from two standard 8-bit units.

Because the transit time of the loop is now one line period, instead of an element period, the speeds of the subtractor, shifter and adder are less important and they can be realised using TTL circuits. In practice it was found convenient to use the same shifter as for the horizontal filter for reasons that will be apparent later.

The increased loop transit-time calls for clocked circuits, that is the subtractor and adder become clocked stages; this in turn alters the transit time of the store. The timing of the loops can be understood by referring to Fig. 14. Not only must the transit time of the subtractor and adder be allowed for but that of the processor must also be included. This last element adds an extra delay to one of the signals reaching the adder, which accounts for the compensating delay in the other adder input. When the total processing time is allowed for, the store is shortened by 5 clock pulses.

6.2. Observed performance

The amplitude/frequency response characteristic was as predicted, i.e. that of Fig. 7 having a maximum attenuation at odd multiples of half line-frequency. The pulse response was also as predicted, i.e. a decaying exponential 'down the picture'. Fig. 15 shows the effect of the filter on Test Card. Note that the frequency bars are unfiltered horizontally. The general effect is that of a vertical leaky integrator.

7. The combination of horizontal and vertical first-order recursive filters

A combined horizontal and vertical first-order recursive filter can be defined through the difference equation

$$y_{ij} = A_{00}x_{ij} - B_{10}y_{i-1,j} - B_{01}y_{i,j-1} - B_{11}y_{i-1,j-1}$$

The corresponding z-transform transfer function is:

$$H(z_1, z_2) = A_{00}/(1 + B_{10}z_1^{-1} + B_{01}z_2^{-1} + B_{11}z_1^{-1}z_2^{-1})$$

From this expression it is clear that A_{00} affects only the gain of the filter and its characteristics are determined by the values of B_{10}, B_{01} and B_{11}.

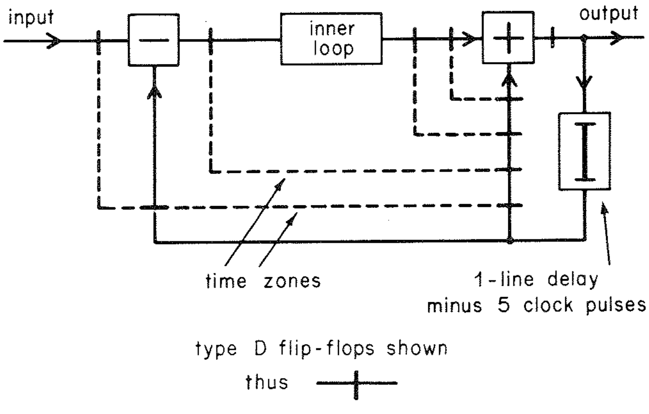
input
inner loop
output
I
time zones
1-line delay minus 5 clock pulses
type D flip-flops shown thus

Fig. 14 - The time partitioning of the vertical filter

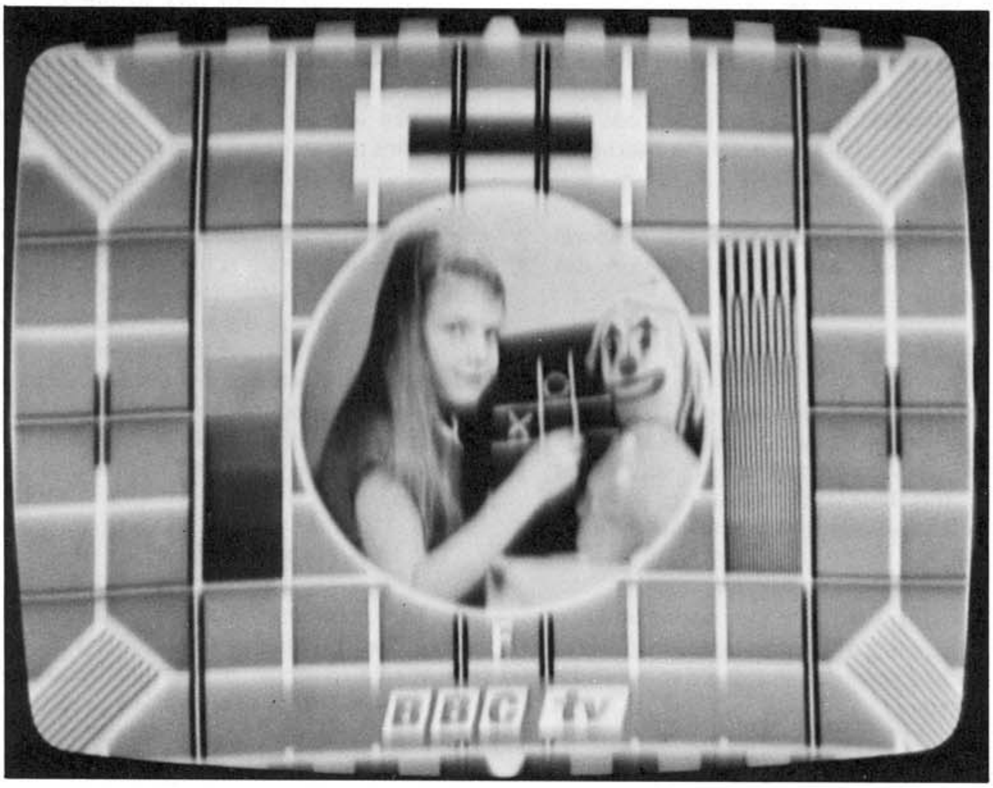

Fig. 15 - The effect of the vertical filter on Test Card

There are obviously many combinations that can exist. Table 3 shows some of the variations. The 'nested' and quasi-parallel forms are particularly interesting because they involve only one parameter and can be easily realised using the basic one-dimensional form developed earlier. Fig. 16 shows their realisations. The cascaded form is also interesting because it is the only one that is variables separable and whose pulse response decreases monotonically in every direction away from the origin; it is thus the nearest to a two-dimensional leaky integrator. However as it involves two parameters it will not be discussed further. Its realisation is also shown in Fig. 16.

Consider first the nested form. The two-dimensional amplitude/frequency and phase/frequency response characteristics are shown in Fig. 17 and the pulse response in

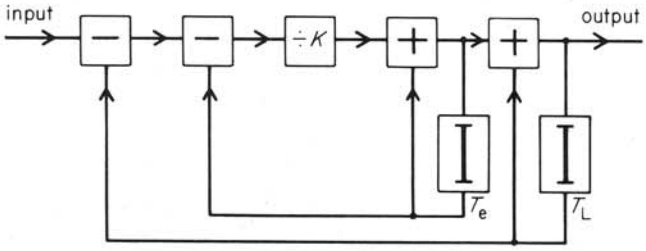

(*a*)

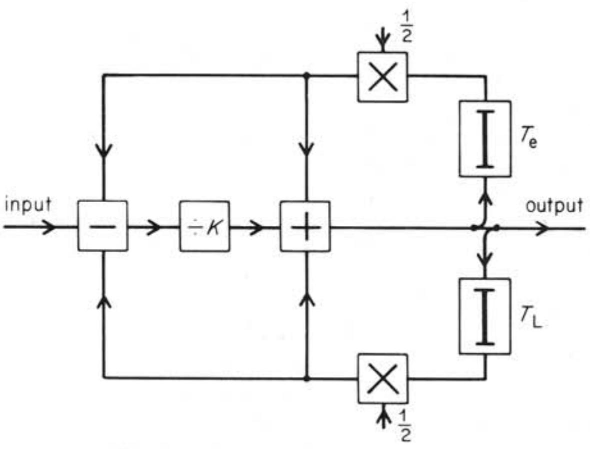

(*b*)

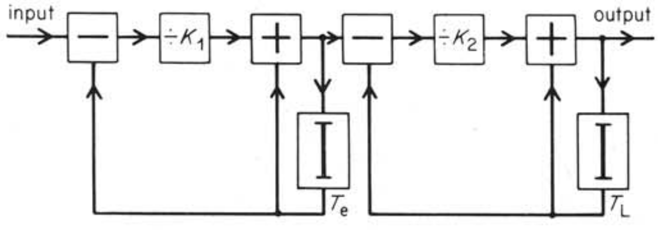

(*c*)

Fig. 16 - Realisation of various two-dimensional filters
(*a*) nested (*b*) quasi-parallel (*c*) cascaded

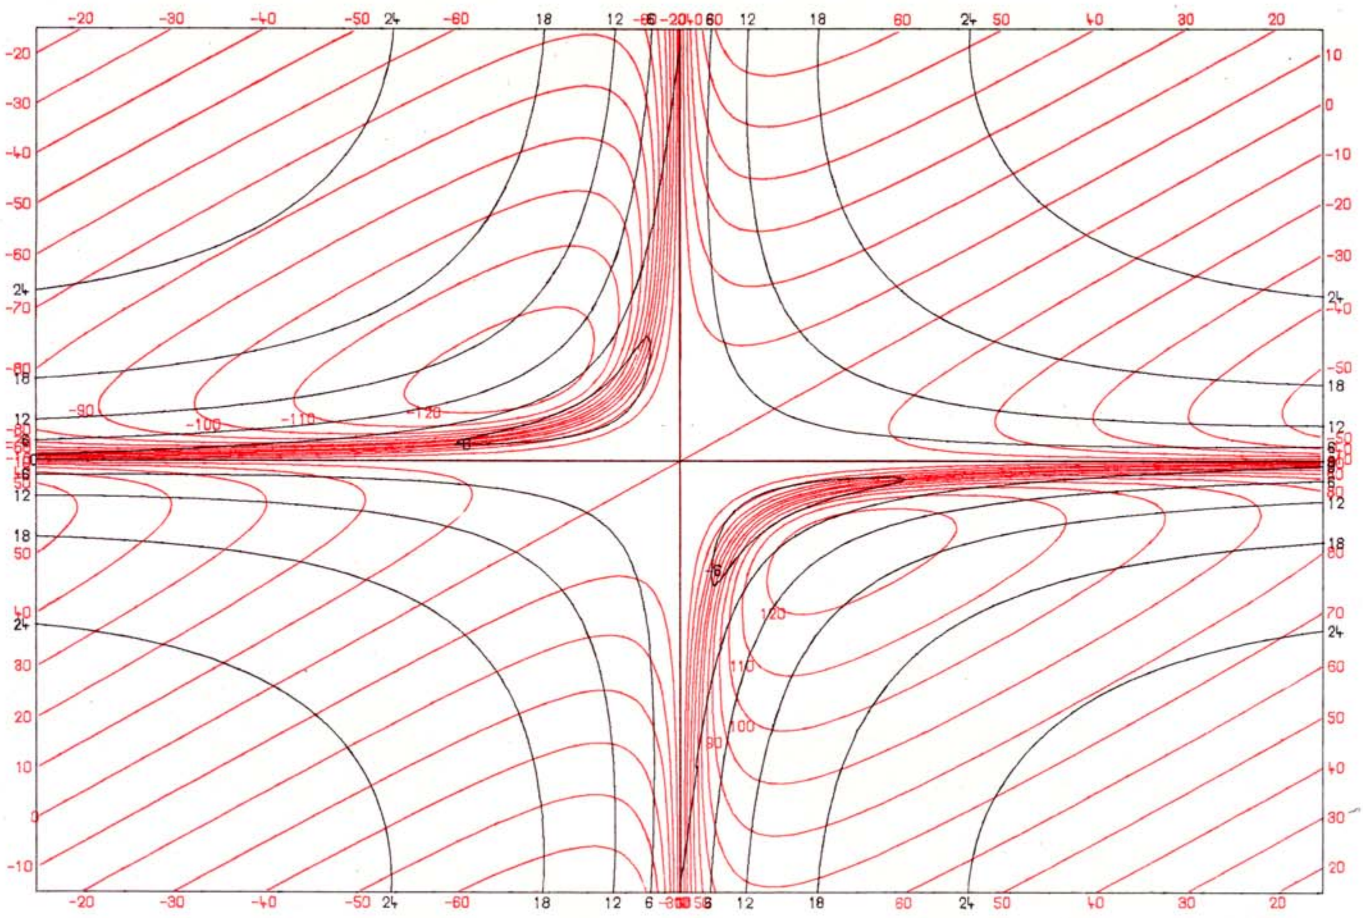

Fig. 17 - Amplitude and phase/frequency characteristics for the nested form

TABLE 3

Various Forms of Two-Dimensional First-Order Recursive Filter

B_{10}	B_{01}	B_{11}	Remarks
b	0	0	pure vertical
0	b	0	pure horizontal
0	0	b	pure skew
b	b	0	quasi-parallel
b	b	$-b$	nested
b	b	b	
b_1	b_2	0	
b_1	b_2	$-b_1 b_2$	cascaded

Fig. 18, for K equal to 8.* In the graph of Fig. 17 the horizontal range is ±5·5 MHz and the vertical range is ±156¼ c/p.h. The axes have been relatively scaled to give the same indication of absolute spatial frequency in c/m.m.

The symmetry of the coefficient pattern assured by $B_{10} = B_{01}$ is reflected in these characteristics. Thus the amplitude and phase characteristics are symmetrical about a skew axis passing through the origin (the deviation of skewness from 45° is caused by the unequal horizontal and vertical sampling intervals).

In the case of the amplitude/frequency response, except for the region near the origin, the contours resemble a family of rectangular hyperbolae; thus the purely horizontal and vertical spatial frequencies are unattenuated. Away from the axes, the attenuation rises rapidly at first, at a rate determined by K. Near the origin, however, in the second and fourth quadrants, there is a trough, and as K increases, this trough deepens and moves nearer the origin. This means that for high values of K the filter has a definite bandpass characteristic so that certain diagonal frequencies in the input are accentuated. This behaviour is very unlike that of the one-dimensional first order filters and means that the filter actually has gain in the passband.

The phase/frequency characteristic shows that a large variation of phase occurs, much larger than the one-dimensional case for the same value of K. The greatest rate of change occurs where the amplitude characteristic is stationary, i.e. at the bottom of the trough. Thus, for the accentuated diagonals the group delay is greatest.

---

* Thanks are due to the staff of mathematics unit, in particular Mr. G. Blakoe, for producing the computer graphics shown in this Section.

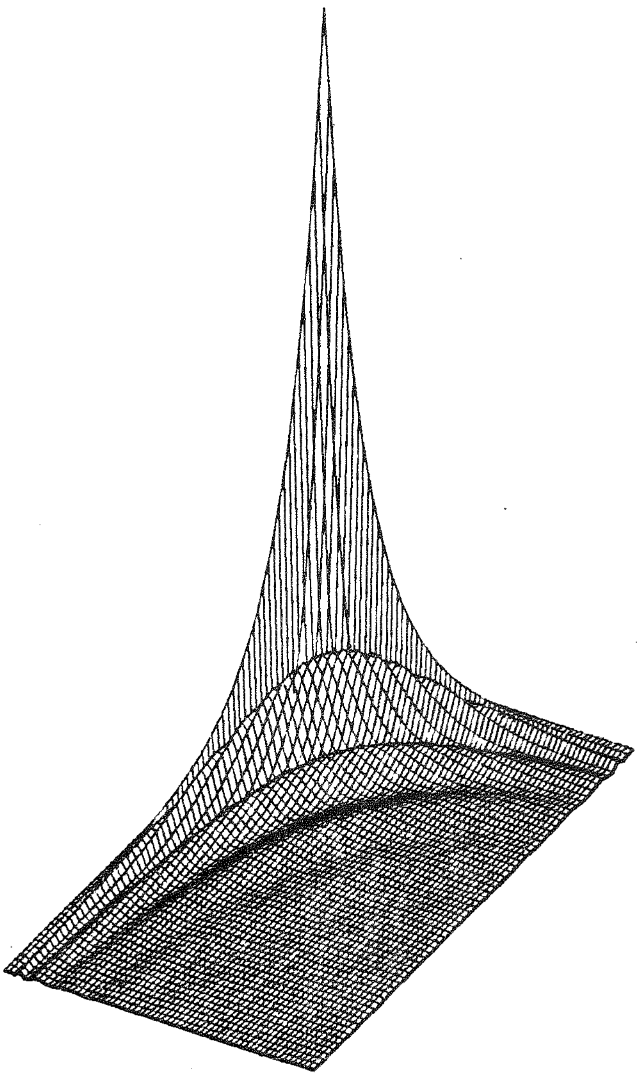

Fig. 18 - Pulse response of nested form

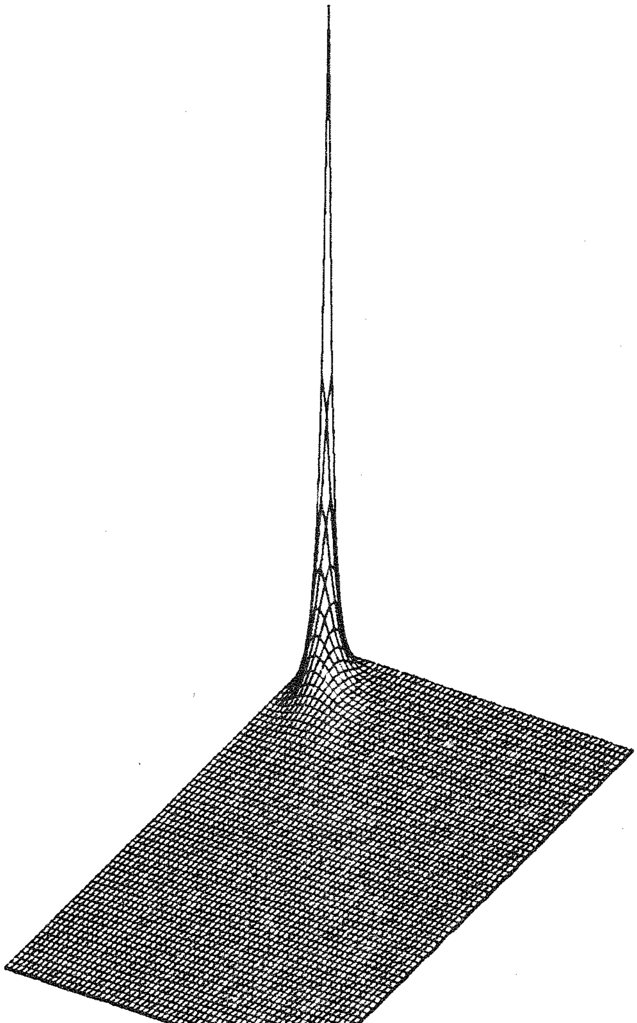

Fig. 19 - Pulse response of quasi-parallel form

In summary the nested form does not, at all, behave like the one-dimensional filters which comprise it, because of the unexpected bandpass behaviour.

Consider now the quasi-parallel form. The two-dimensional frequency characteristics are shown in Fig. 20 and the pulse response in Fig. 19. The characteristics are again symmetrical about the same skew axis as in Fig. 17. But now the amplitude/frequency characteristic is much more gentle with no bandpass effects. In fact the characteristic more nearly approaches that of the ideal isotropic low-pass filter. The phase/frequency characteristic shows that the maximum rate of change of phase again occurs where the amplitude response is stationary, i.e. at the origin. Although the filter is again not phase-corrected, the excursions of phase are much less than for the nested arrangement. It is therefore a more satisfactory arrangement if low-frequency filtering is required.

Fig. 21 shows the effect of the nested form of filter on Test Card. Note that two-dimensional spatial transients i.e. coincident vertical and horizontal changes, cause perturbations similar to the pulse response. This gives rise to a crinkly appearance, the spatial frequency of the perturbations being that at the centre of the passband. The effect is objectionable and shows why the filter cannot be used for low-frequency filtering as it stands.

The effect of the quasi-parallel form of filter on a picture cannot be shown because at the time of writing it has not been constructed. Its realisation is more difficult than that of the nested form because the fast processing loop must be broken to include a further addition.

8. Conclusions

It is perfectly possible to construct first-order recursive horizontal and vertical filters operating on a television signal in 'real time'. The effect of these filters is that of a leaky integrator and they therefore introduce high-frequency loss and are not phase corrected. Their usefulness lies in the fact that it is possible to achieve very great selectivity, i.e. a low cut-off frequency, without the need for many units

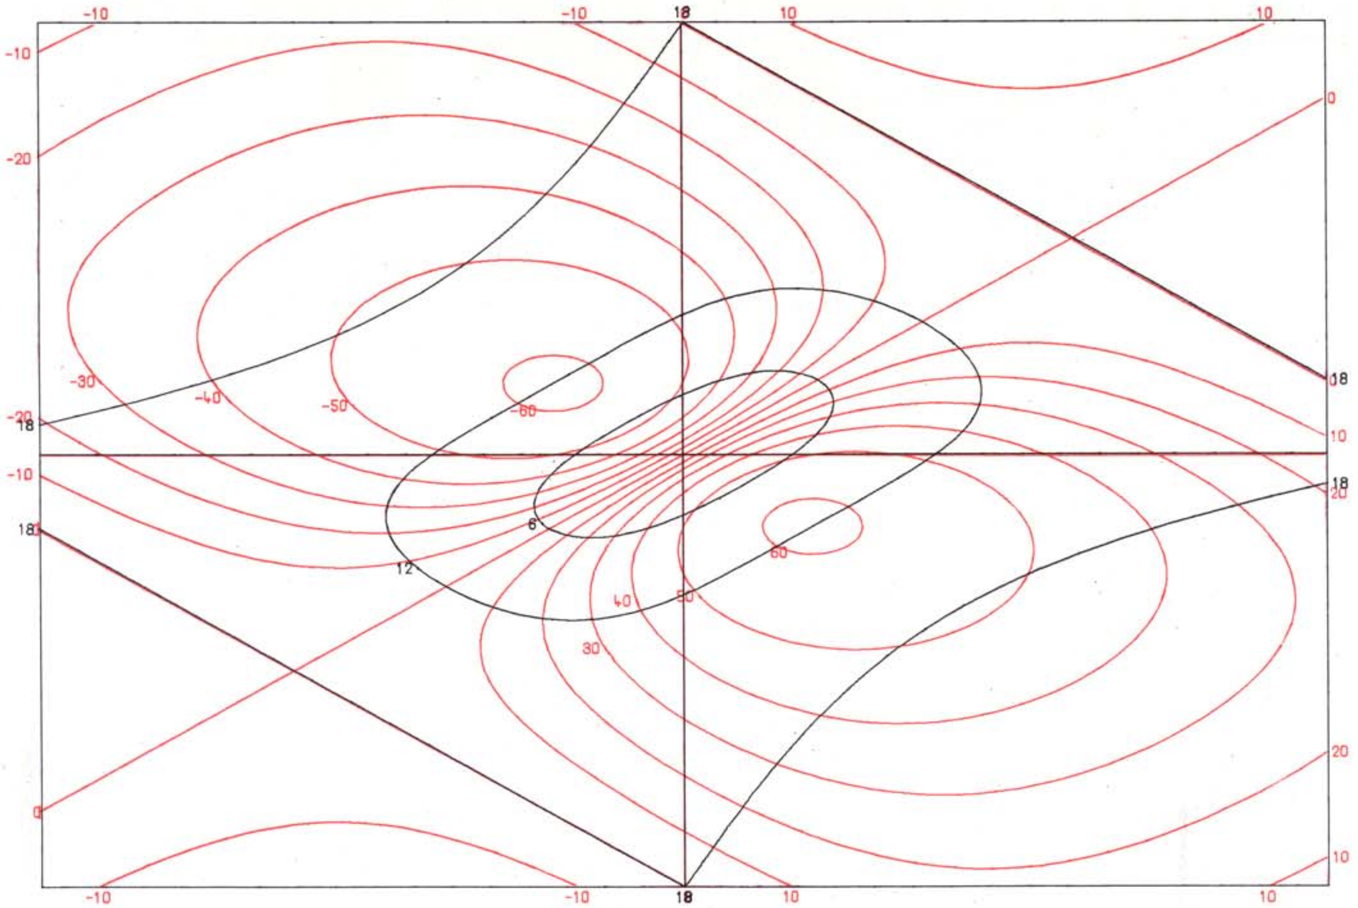

Fig. 20 - Amplitude and phase/frequency characteristics for the quasi-parallel form

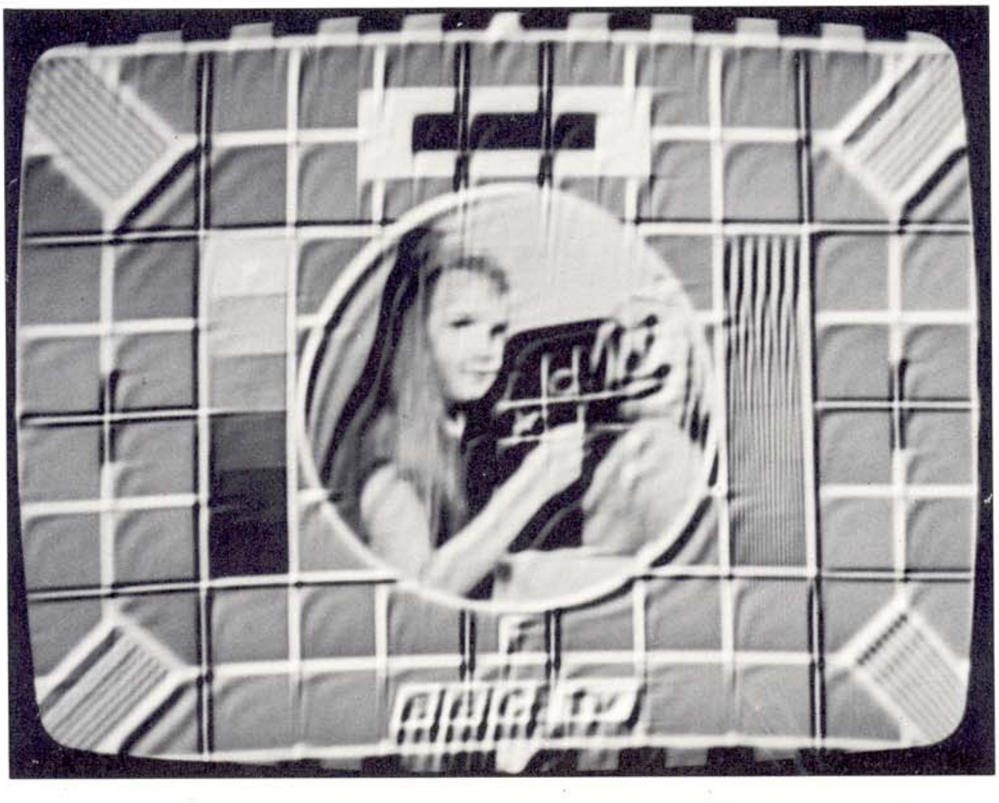

Fig. 21 - The effect of the nested filter on Test Card

of storage as would be the case with non-recursive, i.e. transversal, filters.

It is also possible to combine horizontal and vertical filters to form two-dimensional filters. Two forms of combination have been studied and one form has been built. Such filters offer the possibility of great selectivity in two-dimensional form so that certain spatial frequencies can be accentuated.

This study serves as an introduction to recursive filtering applied to television processing. Although confined to first order filters it has provided much useful information on the problems likely to be encountered with more complex filters. The application of these filters will be discussed further in companion Reports.

9. References

1. See, for example, *IEEE Trans.* on Audio and Electroacoustics Vol. AU–20, No. 4, October 1972.

2. RABINER, L.R. *and* RADER, C.M., editors. Digital signal processing. New York, IEEE Press 1972.

3. GOLD, B. *and* RADER, C.M. Digital processing of signals. New York, McGraw-Hill, 1969.

SMW/JUC

Printed by BBC RESEARCH DEPARTMENT, Kingswood Warren, Tadworth, Surrey, KT20 6NP